arXiv:1609.02502v4 [physics.comp-ph] 29 Dec 2016

Coupling sample paths to the partial thermodynamic limit in stochastic chemical reaction networks

Ethan Levien[1], Paul C. Bressloff[1]

[1]*Department of Mathematics, University of Utah, Salt Lake City, UT 84112 USA*

Abstract

Many biochemical systems appearing in applications have a multiscale structure so that they converge to piecewise deterministic Markov processes in a thermodynamic limit. The statistics of the piecewise deterministic process can be obtained much more efficiently than those of the exact process. We explore the possibility of coupling sample paths of the exact model to the piecewise deterministic process in order to reduce the variance of their difference. We then apply this coupling to reduce the computational complexity of a Monte Carlo estimator. In addition to rigorous results concerning the asymptotic computational complexity of the Monte Carlo estimator, numerical simulations are performed on some simple biological models suggesting that significant computational gains are made.

Key words: chemical reaction networks, Monte Carlo, variance reduction, piecewise deterministic Markov process, Gillespie algorithm

1. Introduction

Large stochastic biochemical reaction networks are a popular modeling framework for investigating cellular processes [1], but the complexity and population sizes involved in realistic models pose major computational challenges. However, when there is a separation of scales, such models lend themselves to a number of model reduction techniques that are useful for course grained analysis. One example occurs when there is a separation in species abundances [2, 3]. If some subset of chemical species in a reaction network are extremely abundant, then reaction channels involving those species will generally occur much faster than reactions involving less abundant species. In this case, one can take a partial thermodynamic limit to obtain a piecewise deterministic Markov process (PDMP). A number of recent studies have provided rigorous errors bounds for this type of reduction [4, 2, 5, 6]. While the PDMP yields useful information about stochastic effects of the rare species, quantitative information about the stochastic fluctuations of the abundant species is lost. On the other hand, in many systems, particularly those with feedback between the rare and abundant chemical species, there is an interest in quantifying the stochastic effects due to these fluctuations [7]. A common method for resolving these fluctuations is the diffusion approximation. While the diffusion approximation is often thought to be computationally advantageous, recent work on classically scaled population models has shown that this methods yields only moderate computational gains [8]. Moreover, the error between the PDMP and the exact model is fixed and it is sometimes desirable to control this quantity, especially when the separation of scales is only moderate.

An alternative to multiscale reduction techniques is to develop methods for accelerating stochastic simulation algorithms such as the Gillespie algorithm [9, 10, 11]. For example, there have been numerous studies on the method of τ-leaping in an effort to accelerate simulations of continuous time Markov chains [12, 13]. More recently, multi-level methods that couple tau-leaping approximations at different resolutions have been

Email addresses: levien@math.utah.edu (Ethan Levien[1]), bressloff@math.utah.edu (Paul C. Bressloff[1])

used to reduce variances in Monte Carlo estimators [14, 15, 8]. Variance reduction techniques that utilize probabilistic couplings have also appeared earlier in the context of SDEs and Markov Chain Monte Carlo methods [16, 17]. While there has been some work that leverages multiscale reduction techniques for Monte Carlo estimators [18], to our knowledge the idea of using these techniques directly as a variance reduction tool has not been studied.

In this paper we explore the idea of coupling reduced models to exact models as a variance reduction tool for Monte Carlo estimators. We develop a new efficient Monte Carlo estimator for multiscale chemical reaction networks near a partial thermodynamic limit. The key insight is that, since only a small fraction of the degrees of freedom of the PDMP are stochastic, one can efficiently compute statistics of the process without Monte Carlo simulations using non Monte Carlo based techniques. On the other hand, if one wants to resolve demographic noise in the full model it is necessary to perform a large number of Monte Carlo simulations. By coupling the full stochastic model to the PDMP one can reduce the variance by a factor inversely related to the system size, and hence a smaller number of simulations need to be performed to achieve a given error tolerance. For practical applications the desired error tolerance of the Monte Carlo estimator scales with this factor. Hence the coupled Monte Carlo estimator has the potential to speed up computations by a fractional power of the error tolerance. Our results extend the idea of variance reduction developed in [14, 15] and provide a new computational application of the theory developed in previous work on PDMP approximations, or partial thermodynamic limits [4, 2, 5].

The paper is organized as follows. In Section 2 we give a introduction to stochastic chemical reaction networks in the single scale, or classical setting. This includes an introduction to the random time change representation of Kurtz, the Gillespie algorithm and Monte Carlo simulations. In Section 3 we introduce multiscale models and show how to obtain a piecewise deterministic Markov process by taking a partial thermodynamic limit. Section 4 introduces a framework for variance reduction in Monte Carlo estimators using the theory in Section 3. Finally, Section 5 and 6 give an analysis of the computational complexity of the coupled Monte Carlo estimator, while numerical examples are presented in Section 7. Non-Monte Carlo based methods for obtaining statistics of the reduced model are presented in Appendix A, while our technical results are proved in Appendix B and C.

2. Representation and simulation of stochastic chemical reaction networks in the classical setting

In this section we briefly introduce some background material pertaining to stochastic chemical reaction networks in the classical setting. Numerous books and articles provide a more comprehensive review of this material [7, 19]. A chemical reaction network provides a mathematical description a set of interacting species, denoted $\mathcal{X} = \{\mathcal{X}_i,\, i = 1, \ldots, d\}$. In general d may be infinite and although our formulation is completely general, we are mostly concerned with examples where d is finite. These interactions are defined by a set of single-step reactions $\mathcal{R} = \{\mathcal{R}_j,\, j = 1, \ldots, p\}$. Let x_i be the number of molecules of $\mathcal{X}_i$ and set $x = (x_1, \ldots, x_d)$. The j-th reaction takes the form

$$\mathcal{R}_j : \sum_{i=1}^{d} k_{j,i}^{\text{in}} \mathcal{X}_i \to \sum_{i=1}^{d} k_{j,i}^{\text{out}} \mathcal{X}_i,$$

where $k_{j,i}^{\text{in}}, k_{j,i}^{\text{out}}$ are known as *stochiometric coefficients*. When such a reaction occurs the state $\mathbf{n}$ is changed according to

$$x_i \to x_i + k_{j,i}, \quad k_{j,i} = k_{j,i}^{\text{out}} - k_{j,i}^{\text{in}}.$$

More complicated multi-step reactions can always be decomposed into these fundamental single-step reactions with appropriate stochiometric coefficients. In practice, most reactions involve collisions between pairs of molecules, so that we take $\sum_i k_{j,i}^{\text{in}} = 1$ or 2.

In the case of a large number of molecules for each species, one often describes the dynamics of a reaction network in terms of a deterministic kinetic or rate equation involving the concentrations $z_j = x_j/\Omega$ – the *law of mass action* [1]. Here Ω is a dimensionless quantity representing the system size, which in gene networks

is typically taken to be the characteristic number of proteins. Alternatively, it could represent some volume scale factor. Currently, we are in the *classical setting* where Ω scales the abundances of every species. This is in contract to the *multiscale setting* where only the size of some subsystem is scaled with the system size.

For a set of p reactions, the kinetic equations take the form (strictly speaking in the thermodynamic limit $\Omega \to \infty$)

$$\frac{dz_i}{dt} = \sum_{j=1}^{p} k_{j,i} \left[\zeta_j \prod_{i=1}^{d} z_i^{k_{j,i}^{\text{in}}} \right] \equiv \sum_{j=1}^{p} k_{j,i} \alpha_j(z), \tag{2.1}$$

where ζ_j is a constant that depends on the probability that a collision of the relevant molecules actually leads to a reaction, and $z = x/\Omega$. The product of concentrations is motivated by the idea that in a well-mixed container there is a spatially uniform distribution of each type of molecule, and the probability of a collision depends on the probability that each of the reactants is in the same local region of space. Ignoring any statistical correlations, the latter is given by the product of the individual concentrations. The functions α_j are known as transition intensities or *propensities*. Given this notation, it is straightforward to write down the corresponding chemical master equation for finite Ω, which takes into account intrinsic fluctuations in the number of molecules (demographic noise). Setting $P(x,t) = \mathbb{P}(x(t) = x | x(0) = x_0)$, the chemical master equation is

$$\frac{dP(x,t)}{dt} = \Omega \sum_{j=1}^{p} \left(\prod_{i=1}^{d} \mathbb{E}^{-k_{j,i}} - 1 \right) \alpha_j(x/\Omega) P(x,t), \tag{2.2}$$

Here $\mathbb{E}^{-k_{j,i}}$ is a step or ladder operator such that for any function $g(x)$,

$$\mathbb{E}^{-k_{j,i}} g(x_1, \ldots, x_i, \ldots, x_d) = g(x_1, \ldots, x_i - k_{j,i}, \ldots, x_d). \tag{2.3}$$

One point to note is that when the number of molecules is sufficiently small, the characteristic form of a propensity function $\alpha(\mathbf{x})$ in equation (2.1) has to be modified:

$$\left(\frac{x_i}{\Omega}\right)^{k_{j,i}^{\text{in}}} \to \frac{1}{\Omega^{k_{j,i}^{\text{in}}}} \frac{x_i!}{(x_i - k_{j,i}^{\text{in}})!}.$$

In this paper, it will be more convenient to consider an alternative representation of a chemical reaction network. If we let the random variable $X_i(t)$ denote the number of molecules in species i, then we can express $X_i(t)$ as

$$X_i(t) = X_i(0) + \sum_j R_j(t) k_{j,i}$$

where $R_j(t)$ denotes the number of times reaction j has occurred by time t. For example, if the propensity function α_j is constant for each j, $R_j(t)$ is a Poisson process with rate $\Omega\alpha_j$, so that letting Π_j denote a unit rate Poisson process gives us

$$X_i(t) = X_i(0) + \sum_j \Pi_j(\Omega \alpha_j t) k_{j,i}.$$

More generally, if α_j depends on $X(t) = \{X_i(t)\}_{i=1,\ldots,d}$ we obtain

$$X_i(t) = X_i(0) + \sum_j \Pi_j \left(\Omega \int_0^t \alpha_j(X(s)/\Omega) ds \right) k_{j,i}. \tag{2.4}$$

This representation is due to Kurtz and is known as the random time change representation of the process $X(t)$ [19]. Unlike the chemical master equation (2.2), which gives a differential equation for the density of $X_i(t)$, the representation (2.4) can be analyzed using probabilistic techniques to obtain meaningful information, even for extremely complicated reaction networks. Additionally, since it is a pathwise representation rather than a description of a density it is more useful as a basis for deriving simulation algorithms. While one usually thinks of (2.2) or (2.4) as representing a well-mixed chemical reaction, this framework is easily

extended to spatial models of reaction diffusion systems where the diffusion is represented by additional reaction channels [5]. Hence, although we will usually speak of non-spatial reaction networks in this paper, our results can be extended to spatial models.

There are two main methods for generating exact realizations of the paths (2.4). These are the Next Reaction method, and the Gillespie algorithm. We will present only the Gillespie algorithm in, which is given in Algorithm 1, since it is more easily related to other algorithms used in this paper. In particular it can be modified to simulate piecewise deterministic Markov process in a fairly straightforward manner. Let us denote

$$\mathsf{C}_X := \mathbb{E}[\ \#\text{ of computations to simulate } X(T) \text{ using Algorithm 1}].$$

More generally, we define the *complexity* of an algorithm as the order of magnitude of the expected simulation time with respect to some parameter in the model. For our proposes this will be the system size Ω. To get a handle on C_X, note that $\mathbb{E}[t_{\text{next}}] = O(\Omega^{-1})$ implying we expect to simulate $T\mathbb{E}[t_{\text{next}}]^{-1} = O(\Omega)$ events, so that $\mathsf{C}_X = O(\Omega)$. It should be noted that the most computationally expensive step in each iteration of Algorithm 1 is the computation of the reaction index k, and the next reaction method can be modified to reduce the cost of this step [20]. On the other hand, it is the $O(\Omega^{-1})$ scaling of the complexity that is important for our analysis.

Algorithm 1 Simulation of (2.4)

1: Initialize $X(0)$ and set $t = 0$.
2: **while** $t \le T$ **do**
3: $\quad \alpha_{\text{total}} := \Omega \sum_j \alpha_j(X(s)/\Omega)$
4: $\quad$ Generate random numbers $t_{\text{next}} \sim \text{Exp}(\alpha_{\text{total}})$ and $r \sim \text{Unif}(0,1)$
5: $\quad \bar{j} := \min_i\{i : \sum_{j=1}^{i} \alpha_j(X(s)) < r\alpha_{\text{total}}\}$
6: $\quad X(t + t_{\text{next}}) := X(t) + k_{\bar{j}}$
7: $\quad t := t + t_{\text{next}}$

2.1. Monte Carlo simulations in the classical setting

Suppose we wish to approximate some statistics of some process $X_k(T)$ using a crude Monte Carlo estimator (MCE) $\widehat{Q}_{\text{crude}}(M)$ with M realizations of the process. That is,

$$\widehat{Q}_{\text{crude}}(M) = \frac{1}{M}\sum_{j=1}^{M} f(X_{k,[j]}(T)). \tag{2.5}$$

Here and elsewhere we use the convention that the subscript $[j]$ indicates a specific realization of a process. In order for $\widehat{Q}_{\text{crude}}(M)$ to approximate $\mathbb{E}[f(X_k(T))]$ to order ε in the sense of confidence intervals, that is, in the sense that the standard deviation of $\widehat{Q}_{\text{crude}}(M)$ is order ε, we need

$$\text{Var}(\widehat{Q}_{\text{crude}}(M)) = O(\varepsilon^2)$$

and hence $M = O(\varepsilon^{-2}\text{Var}(f(X_{k,[j]}(T)))$. It is important to note that Ω and ε must be related. To see how such a relationship comes about, observe that in practice ε will scale with Ω since we need to decrease ε to resolve demographic noise as Ω increases. Following [8], we will set

$$\varepsilon = \Omega^{-\delta}$$

where $\delta \ge 0$ is a measure of the accuracy of the Monte Carlo estimator relative to the system's noise. Then the complexity of the MCE $\widehat{Q}_{\text{crude}}(M)$ is

$$\mathsf{C}_{\text{crude}} = O(\varepsilon^{-2}\mathsf{C}_X\text{Var}(f(X_k(T))).$$

It can be shown that $\mathrm{Var}(f(X_k(T)) = O(\Omega^{-1})$, which corresponds to the stochastic model approaching a deterministic limit at a rate inversely proportionally to the size of the system. On the other hand, $\mathsf{C}_X = O(\Omega)$, so that in the classical setting, the contribution of the complexity from the simulation of the path cancels with the variance of the path, and we obtain

$$\mathsf{C}_{\text{crude}} = O(\varepsilon^{-2}).$$

Essentially, when a stochastic model approaches a deterministic limit, variance reduction in terms of the system size is "for free". This applies not only to the crude MCE, but any Monte Carlo method applied to a biochemical model in the classical setting. We refer to [8] for a rigorous analysis of Monte Carlo methods in the classical setting. The crucial observation that motivates the developments in this paper is that in the multiscale setting, where some species abundances do not scale with Ω, it is generally not possible to bound $\mathrm{Var}(f(X_k(T))$ in terms of Ω, and the complexity of any Monte Carlo method increases by an order of magnitude.

3. Multiscale reaction networks and the partial thermodynamic limit

We want to consider multiscale chemical reaction networks involving $l + m$ species, where there are l abundant species $\{\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_l\}$ and m rare species $\{\mathcal{Y}_1, \mathcal{Y}_2, \ldots, \mathcal{Y}_m\}$. We will take X_i and Y_i to represent the number of species $\mathcal{X}_i$ and $\mathcal{Y}_i$ respectively. The motivation behind this notation is that when $m = 0$ we retrieve the classical setting both notationally and conceptually. In some applications one refers to $X = (X_i)_{i=1,\ldots,l} \in \mathbb{N}^l$ as the population and $Y = (Y_i)_{i=1,\ldots,m} \in \mathbb{N}^m$ as the environment [21], and we will often adopt this terminology. In order to justify the labelings of rare and abundant species we need two main assumptions. First, we assume that only the initial concentration of the abundant species scales with the systems size.

Assumption 1 (separation of scales). *Let Ω denote the system-size introduced in Section 2. Then*

$$||X(0)||_1 = O(\Omega), \quad ||Y(0)||_1 = O(1)$$

Here and throughout the rest of this paper the $O(\cdot)$ terms are with respect to Ω. Note that for any vector $x \in \mathbb{R}^l$, $||x||_1 = O(g(\Omega))$ is equivalent to saying $||x||_\infty = \max_j |x_j| = O(g(\Omega))$ so that this assumptions implies $|X_i(0)| = O(\Omega)$ for each $i = 1, \ldots, l$.

Physically, Ω can be thought to control the volume $V(\Omega)$ of a well-mixed domain such that the density of each abundant species and the number of each rare species (environment) are independent of the volume. We will assume that $V(\Omega)$ is linearly increasing in Ω.

Following the notation of section 2, each reaction in $\mathcal{R} = \{\mathcal{R}_j,\ j = 1, \ldots, p\}$ has the form

$$\mathcal{R}_j : \sum_{i=1}^{m} k_{j,i}^{\text{in}} \mathcal{X}_i + \sum_{i=1}^{l} u_{j,i}^{\text{in}} \mathcal{Y}_i \to \sum_{i=1}^{m} k_{j,i}^{\text{out}} \mathcal{X}_i + \sum_{i=1}^{l} u_{j,i}^{\text{out}} \mathcal{Y}_i \tag{3.1}$$

The reaction vectors are given by

$$k_j = (k_{j,1}^{\text{out}} - k_{j,1}^{\text{in}}, \ldots, k_{j,l}^{\text{out}} - k_{j,l}^{\text{in}})^T \in \mathbb{Z}^l$$
$$u_j = (u_{j,1}^{\text{out}} - u_{j,1}^{\text{in}}, \ldots, u_{j,m}^{\text{out}} - u_{j,m}^{\text{in}})^T \in \mathbb{Z}^m.$$

Suppose that only a proper subset of the reactions involve changes in the state of the rare species and define

$$\mathcal{J}_1 = \{j : u_{j,i'} \neq 0, \text{ for some } i' \in \{1, \ldots, m\}\}, \quad \mathcal{J}_0 = \{1, \ldots, p\} - \mathcal{J}_1.$$

That is, $\{R_j : j \in \mathcal{J}_1\}$ is the set of reactions that produce or annihilate at least one of the rare species. As a further simplification, we assume that this set of reactions does not produce changes in the number of any

abundant species. The random time change representation of networks of this form can then be written in the form

$$X_i(t) = X_i(0) + \sum_{j \in \mathcal{J}_0} \Pi_j \left(\int_0^t \Omega \alpha_j(X(s)/\Omega, Y(s)) ds \right) k_{j,i}, \quad i = 1, \ldots, l \tag{3.2a}$$

$$Y_{i'}(t) = Y_{i'}(0) + \sum_{j \in \mathcal{J}_1} \Pi_j \left(\int_0^t \alpha_j(X(s)/\Omega, Y(s)) ds \right) u_{j,i'}, \quad i' = 1, \ldots, m. \tag{3.2b}$$

where Π_j are once again independent unit rate Poisson processes. Note that only the number of each abundant species is scaled by the system size Ω. We will always assume that the rate functions $\alpha_j(x, y)$ are polynomials in x and Lipschitz continuous x uniformly in y, that is, there exists $L_{\alpha_j}(K)$ such that for all $||x_1 - x_2|| < K$,

$$|\alpha_j(x_1, y) - \alpha_j(x_2, y)| \leq L_{\alpha_j}(K)|x_1 - x_2| \tag{3.3}$$

for all $y \in \mathbb{N}^m$. This is consistent with the law of mass action for the abundant species, that is,

$$\alpha_j(x, y) = \zeta_j(y) \prod_{i=1}^{m} x_i^{k_{j,i}^{\text{in}}}. \tag{3.4}$$

We will assume $\zeta_j(y) = O(1)$ with respect to all other parameters, although with a more sophisticated scaling than the one carried out below our analysis can be extended to the case where ζ_j vary over many orders of magnitude [15]. Scaled population models of this sort appear in a number of applications, such as gene networks where the promoters have a fixed number of states but the abundance of a gene scales with other properties (such as the size) of the cell [22]. We would hope that the scaled rates along with suitable restrictions on the dynamics ensure the initial separation of scales is preserved over bounded time intervals. This usually involves some notion of stability, or regularity, for the stochastic equations (3.2) and rigorous results concerning the stochastic stability of continuous time Markov chains can established for general biochemical reaction models [5]. The focus of this paper is not stochastic stability, but rather variance reduction in Monte Carlo simulations, so to keep our presentation self contained we will make the following assumption.

Assumption 2 (stability). *Let $(X(t), Y(t))$ satisfy equation (3.2). Then $||Y(t)||_1 \leq c$ for all $t < T$ and for all $\sigma \geq 1$,*

$$\mathbb{E}[||X(t)||_1^\sigma] < ||X(0)||_1^\sigma B_\sigma(T) \tag{3.5}$$

where $B_\sigma(t)$ is independent of $||X(0)||_1$. Moreover, if we define the stopping time

$$\tau_K = \inf_{t \geq 0} \{t : ||X(t)||_1 \geq K||X(0)||_1\}, \tag{3.6}$$

then for all $T > 0$ there exists a positive constant $C(T)$ independent of K and the initial conditions such that

$$\mathbb{P}(\tau_K < T) \leq C(T)K^{-1}. \tag{3.7}$$

Note that the first bound simply says that the growth rates of the moments are not dependent on the system size. Now let us briefly explain the intuition behind the last statement in Assumption 2, which is less transparent. We begin by noting that the most desirable condition would be the existence of K such that $||X(t)||_1 \leq K||X(0)||_1$ for all $t \leq 0$. This ensures $||X(0)||_1 = O(\Omega)$ and $||Y(0)||_1 = O(1)$, which is exactly what we would like to achieve. For a deterministic system this a sensible condition, but it is far too strong a restriction on the stochastic process $X(t)$. In fact, we can only guarantee sample paths of (3.2) will be bounded with probability one when there is a conservation law in effect ($||X(t)||_1$ and $||Y(t)||_1$ are constant). This suggests we must settle for a more realistic assumption, one that quantifies the event of a sample path breaking these inequalities and ensures K can be chosen to make these events rare. We achieve this in Assumption 2 by bounding the probability that $\tau_K > T$. If the corresponding mean field process

derived below has a trapping region we would naturally expect Assumption 2 to hold since the deterministic dynamics arising in the limit $\Omega \to \infty$ are relatively bounded. More generally, assumption 2 holds for any model with only linear growth, see Example 3.3.

Under these assumptions, we can introduce the scaled population variables, or densities, $X^\Omega(t) = \Omega^{-1}X(t)$, ensuring that $||X^\Omega(t)||_1 = O(1)$ in some bounded interval. The scaling of the rates leads to the scaled dynamics

$$X_i^\Omega(t) = X_i^\Omega(0) + \sum_{j\in\mathcal{J}_0} \Pi_j \left(\Omega \int_0^t \alpha_j(X^\Omega(s), Y(s))ds\right) \Omega^{-1}k_{j,i}, \quad i = 1,\dots,l \tag{3.8a}$$

$$Y_{i'}(t) = Y_{i'}(0) + \sum_{j\in\mathcal{J}_1} \Pi_j \left(\int_0^t \alpha_j(X^\Omega(s), Y(s))ds\right) u_{j,i'}, \quad i' = 1,\dots,m. \tag{3.8b}$$

In light of this scaling, it makes sense to take the thermodynamic limit $\Omega \to \infty$ only in the abundant species, following [4] we call this the *partial thermodynamic limit* (PTDL). Setting $\lim_{\Omega\to\infty} X^\Omega(t) = Z(t)$ in some sense we would expect that $Z_i(t)$ satisfies (3.8a) with the stochastic integral for the abundant species by mean field dynamics. That is,

$$Z_i(t) = Z_i(0) + \sum_{j\in\mathcal{J}_0} \int_0^t \alpha_j(Z(s), Y(s))k_{j,i,}, \quad i = 1,\dots,l \tag{3.9a}$$

$$Y_{i'}(t) = Y_{i'}(0) + \sum_{j\in\mathcal{J}_1} \Pi_j \left(\int_0^t \alpha_j(Z(s), Y(s))ds\right) u_{j,i'}, \quad i' = 1,\dots,m. \tag{3.9b}$$

Another way to express the dynamics of $Z_i(t)$ is through the differential equation

$$Z_i'(t) = \sum_{j\in\mathcal{J}_0} \alpha_j(Z(s), Y(s))k_{j,i}, \quad i = 1,\dots,l \tag{3.10}$$

which holds between discrete jumps. $Z(t)$ is an example of a piecewise deterministic Markov process (PDMP). A comprehensive review of such process along with more rigorous constructions can be found in [23]. While there are many technical statements one can make about various notions of the convergence $X^\Omega(t) \to Z(t)$, the technical results necessary for our analysis are contained in Sections 5 and 6.

The PDMP (3.10) should be compared to the piecewise SDE one would obtain if instead of taking the fully deterministic mean field approximation of $X^\Omega(t)$, one carried out a system-size expansion [24, 1] to leading order. This piecewise Itô SDE is given by

$$dZ_i(t) = \sum_{j\in\mathcal{J}_0} \left(\alpha_j(Z(s), Y(s))dt + \Omega^{-1/2}\sqrt{\alpha_j(Z(s), Y(s))}dW_j(t)\right) k_{ji} \tag{3.11}$$

where $W_j(t)$ are independent Brownian motions. Previous studies have developed hybrid frameworks for simulating multi-scale chemical reaction networks using this system-size expansion [25, 18]. However, in contrast to the Monte Carlo methods discussed in the following sections, these hybrid methods have an error which is fixed in terms of Ω.

3.1. Generating sample paths of the PDMP

Unlike the jump process (3.8), exact realizations of $Z(t)$ cannot be generated due to the continuous nature of (3.10). Instead one must pick an accuracy h and compute an approximation $(Z^h(t), Y^h(t))$ using standard ODE methods between jumps, and then incorporate the stochastic jumps in some suitable way. Note that one only needs to include the stochastic contribution of reaction channels indexed by the set $\mathcal{J}_1$. If the model has no feedback than one can compute $Y^h(T) = Y(T)$ using an exact algorithm, and then solve for $Z^h(T)$ using a suitable ODE method between jumps. For models with feedback, the computation of $Z^h(T)$

is more subtle, since one needs to compute the jumps times while evolving the process $Z^h(T)$. A number of algorithms exist for accomplishing this task. Popular among those studying PDMP approximations of chemical reaction networks are splitting methods, which can be derived by applying operator splitting to the forward Kolmogorov equation for the PDMP [5]. Like τ-leaping, splitting methods obtain approximate jumps times by holding the jump rates constant in $Z(t)$ and evolving $Y(t)$ over some small time interval. We have found a more direct approach, where the next jump time is computed exactly, to be more effective for the models considered in this paper. We leave a more systematic comparison of methods for solving the PDMP arising in thermodynamic limit to future studies.

We now describe direct method. Given the state of the system, $(Z(t_0), Y(t_0)) = (z_0, y_0)$, at time t_0, it can be shown that the next jump time is given by

$$t_1 = \inf\left\{u > t_0 : \int_{t_0}^{u} \sum_{j\in\mathcal{J}_2} \alpha_j(\Phi^s_{y_0}(z_0), y_0)ds = \bar{s}\right\} \tag{3.12}$$

where $\bar{s} \sim \text{Exp}(1)$ and Φ^t_y is the continuous flow of $Z(t)$ when $Y(t) = y$ [26]. In general, the process of solving for $(Z^h(t), Y^h(t))$ at the jump times by numerical solving (3.12) is known as the True Jump Method. One can show that the solution to the minimization problem (3.12) along with the state of the continuous variable at t_1 is given by $(Z(t_1^-), t_1) = (z(\bar{s}), \tau(\bar{s}))$ where

$$\begin{cases} z'(s) = \sum_{j\in\mathcal{J}_1} \alpha_j(z(s), y_0) \left(\sum_{j\in\mathcal{J}_2} \alpha_j(z(s), y_0)\right)^{-1} \\ \tau'(s) = \left(\sum_{j\in\mathcal{J}_1} \alpha_j(z(s), y_0)\right)^{-1} \\ z(0) = z_0, \quad \tau(0) = t_0. \end{cases} \tag{3.13}$$

The process of solving the PDMP using by (3.12) to obtain the solutions to (3.13) has been coined the CHV method [27]. One subtly is that a direct implementation of the CHV method gives only the values of $Z(t)$ at the jump times, whereas we will be interested in obtaining the value of the process at some specified time T. However, this can easily be achieved by integrating the continuous component of the process from the last jump time before T, up to T. The details of the CHV method are given in Algorithm 2. At leading order in h, the complexity of obtaining $Z^h(T)$ with this algorithm is simply the complexity using whatever integration scheme we use to solve (3.13) up to time T. In particular, using high order methods renders the cost of computing $Z^h(T)$ negligible relative to the cost of computing $X(T)$ exactly when Ω is large.

Algorithm 2 Simulation of (3.9)

Select an accuracy h. Initialize $Z^h(0)$ and set $t = 0$.
while $t \leq T$ **do**
  Generate a random number $\bar{s} \sim \text{Exp}(1)$.
  Let $(z(\bar{s}), \tau(\bar{s}))$ be the solution to (3.13) with $t_0 = t$ and $z_0 = Z^h(t)$.
  $Z^h(t + \tau(\bar{s})) = z(\bar{s})$
  $\alpha_{\text{total}} = \sum_{j\in\mathcal{J}_1} \alpha_j(Z^h(\tau(\bar{s})), Y(t))$
  Generate a random number $r \sim \text{Unif}(0, 1)$
  $\bar{j} := \min_i\{i : \sum_{j\in\mathcal{J}_2: j<i} \alpha_j(Z^h(s), Y^h(s)) < r\alpha_{\text{total}}\}$
  $Y^h(\tau(\bar{s})) := Y^h(t) + u_{\bar{j}}$
  $t := \tau(\bar{s})$
Perform numerical integration to obtain $Z^h(T)$.

3.2. Example: Two State Catalyst

Our first example of a simple chemical reaction network without feedback consists of four species $\mathcal{X}_1, \mathcal{X}_2, \mathcal{Y}_1$ and $\mathcal{Y}_2$ obeying the reaction scheme

$$\begin{aligned} \mathcal{X}_1 + \mathcal{Y}_1 &\underset{\alpha}{\to} \mathcal{X}_2 + \mathcal{Y}_1 \\ \mathcal{X}_2 + \mathcal{Y}_2 &\underset{\beta}{\to} \mathcal{X}_1 + \mathcal{Y}_2 \\ \mathcal{Y}_1 &\underset{\lambda}{\to} \mathcal{Y}_2 \\ \mathcal{Y}_2 &\underset{\mu}{\to} \mathcal{Y}_1 \end{aligned} \tag{3.14}$$

where $X_1(0) + X_2(0) = \Omega \gg 1$ and $Y_1 + Y_2 = 1$. This reaction network can model a catalytic reaction [5], or a population of particles jumping between two lattice sites with gate controlled by the Y species [21]. In the former context Y_i would be the catalysts, while in the latter they would represent the state of the barrier between two compartments and X_i would represent the number of particles in each compartment.

In terms of Kurtz's time change representation,

$$\begin{aligned} X(t) &= X(0) - \Pi_1\left(\int_0^t X(s)Y(s)\alpha ds\right) + \Pi_2\left(\int_0^t (\Omega - X(s))(c - Y(s))\beta ds\right) \\ Y(t) &= Y(0) - \Pi_3\left(\int_0^t \lambda(c - Y(s))ds\right) + \Pi_4\left(\int_0^t \mu Y(s)ds\right) \end{aligned}$$

where we have used the conservation laws $X_2(t) = \Omega - X_1(t)$, $Y_2(t) = c - Y_1(t)$ and set $X_1(t) = X(t)$, $Y_1(t) = Y(t)$. The scaled model is

$$\begin{aligned} X^\Omega(t) &= X^\Omega(0) - \Pi_1\left(\Omega\int_0^t X^\Omega(s)Y(s)\alpha ds\right)\Omega^{-1} \\ &\quad + \Pi_2\left(\Omega\int_0^t (1 - X^\Omega(s))(c - Y(s))\beta ds\right)\Omega^{-1} \\ Y(t) &= Y(0) - \Pi_3\left(\int_0^t \lambda(c - Y(s))ds\right) + \Pi_4\left(\int_0^t \mu Y(s)ds\right) \end{aligned}$$

In the limit $\Omega \to \infty$ we obtain the PDMP

$$\begin{aligned} Z(t) &= Z(0) - \int_0^t Z(s)Y(s)\alpha ds + \int_0^t (\Omega - Z(s))(c - Y(s))\beta ds \\ Y(t) &= Y(0) - \Pi_3\left(\int_0^t \lambda(c - Y(s))ds\right) + \Pi_4\left(\int_0^t \mu Y(s)ds\right) \end{aligned}$$

Note that assumption 2 holds due to the conservation laws.

3.3. Example: Switching Gene

A second example of a network without feedback is a simplified model of protein production in the presence of a switching gene [22, 28]:

$$\begin{aligned} \mathcal{Y}_1 &\underset{\Omega\alpha}{\to} \mathcal{Y}_1 + \mathcal{X} \\ \mathcal{X} &\underset{\beta}{\to} \emptyset \\ \mathcal{Y}_1 &\underset{\lambda}{\to} \mathcal{Y}_2 \\ \mathcal{Y}_2 &\underset{\mu}{\to} \mathcal{Y}_1 \end{aligned}$$

In this reaction network $\mathcal{X}$ represents a protein and $\mathcal{Y}_i$ the state associated gene. When the gene is in state $\mathcal{Y}_1$, the protein is produced at a rate proportional to the system size Ω, while when in state $\mathcal{Y}_2$ the protein is not produced. For example, $\mathcal{Y}_1$ could represent the presence of a repressor occupying the RNA polymerase binding site, thereby preventing the transcription of the gene [29, 30]. In this simplified model the complex mechanisms of transcription and translations are viewed as a "back box" represented by the production rate α. The stochastic equations are

$$\begin{aligned} X(t) &= X(0) + \Pi_1\left(\int_0^t Y(s)\Omega\alpha ds\right) - \Pi_2\left(\int_0^t X(s)\beta ds\right) \\ Y(t) &= Y(0) - \Pi_3\left(\int_0^t \lambda(1-Y(s))ds\right) + \Pi_4\left(\int_0^t \mu Y(s)ds\right) \end{aligned}$$

where we have once again used conservation to obtain a two dimensional system. In the partial thermodynamic limit we obtain

$$\begin{aligned} Z(t) &= Z(0) + \int_0^t Y(s)\alpha ds - \int_0^t Z(s)\beta ds \\ Y(t) &= Y(0) - \Pi_3\left(\int_0^t \lambda(1-Y(s))ds\right) + \Pi_4\left(\int_0^t \mu Y(s)ds\right) \end{aligned}$$

In order to establish that Assumption 2 one only needs to note that this assumption holds for a Poisson process which dominates $X(t)$.

3.4. Example: Switching Gene with Feedback

A more complex model of gene expression might include feedback between the protein associated with the gene and environment controlling the gene expression [22]. For example, suppose protein $\mathcal{X}$ activates the repressor blocking the RNA polymerase. This leads to the reaction network

$$\begin{aligned} \mathcal{Y}_1 &\underset{\Omega\alpha}{\to} \mathcal{Y}_1 + \mathcal{X} \\ \mathcal{X} &\underset{\beta}{\to} \emptyset \\ \mathcal{X} + \mathcal{Y}_1 &\underset{\Omega^{-1}\lambda}{\to} \mathcal{X} + \mathcal{Y}_2 \\ \mathcal{Y}_2 &\underset{\mu}{\to} \mathcal{Y}_1 \end{aligned}$$

The stochastic equations are

$$\begin{aligned} X(t) &= X(0) + \Pi_1\left(\int_0^t Y_1(s)\Omega\alpha ds\right) - \Pi_2\left(\int_0^t X(s)\beta ds\right) \\ Y(t) &= Y(0) - \Pi_3\left(\int_0^t \Omega^{-1}\lambda X(s)(1-Y(s))ds\right) + \Pi_4\left(\int_0^t \mu Y(s)ds\right) \end{aligned}$$

The partial thermodynamic limit is

$$\begin{aligned} Z(t) &= Z(0) + \int_0^t Y(s)\alpha ds - \int_0^t Z(s)\beta ds \\ Y(t) &= Y(0) - \Pi_3\left(\int_0^t Z(t)\lambda(1-Y(s))ds\right) + \Pi_4\left(\int_0^t \mu Y(s)ds\right). \end{aligned}$$

Since the $\mathcal{J}_1$ channels have rates that depend on the abundant variables, we have a network with feedback. An argument similar to the one made for example 3.3 implies that assumption 2 holds.

4. Variance reduction in the Multiscaling setting

As noted at the end of Section 2, for multiscale models we can generally not bound the sample variance in terms of the system size and the asymptotic complexity of Monte Carlo methods picks up a factor of Ω. For example,

$$C_{\text{crude}} = O(\varepsilon^{-2-1/\delta}).$$

Loosely speaking, while a great deal of information about the exact model is still contained in the thermodynamic limit $Z(t)$, it is not being used in the computations. Again, we emphasis how this contrasts the classical setting where information about the deterministic limit is used to accelerate the converge of an MCE without any additional work. Generally speaking, our goal is to understand how information about the thermodynamic limit can be used in the multiscale setting.

The coupled Monte Carlo estimator we will introduce is based on the idea of variance reduction via a probabilistic coupling of the exact process with an approximate process. In our case the approximate process will be the PDMP $Z(t)$. This idea has proven to be very useful in, and is the basis for, multilevel Monte Carlo methods [15, 17] where different τ-leaping approximations are coupled. To construct an MCE in the present setting we note that

$$\mathbb{E}[f(X_i(T))] = \mathbb{E}[f(X_i(T)) - f(Z_i(T))] + \mathbb{E}[f(Z_i(T))] \tag{4.1}$$

Two observations allow us to use this decomposition to obtain statistics of $\mathbb{E}[f(X_i(T))]$ more efficiently than the crude MCE (2.5). First, statistics of $Z(T)$ can be obtained much more efficiently than statistics of the exact process when Ω is even moderately large relative to the abundance of the rare species. These statistics can be obtained either by a MCE using Algorithm 2 to generate the sample paths, or by non-Monte Carlo based methods given in Appendix A, which are difficult to apply to the exact process. Second, we can reduce the number of simulations we need to perform of the full process by coupling the processes $X(t)$ and $Z(t)$ in a way that reduces the variance of the difference $f(X_i(T)) - f(Z_i(T))$.

For the second term in (4.1) let us assume we can produce an approximation $\widehat{Q}_Z(h) \approx \mathbb{E}[f(Z_i(T))]$ satisfying

$$|\widehat{Q}_Z(h) - \mathbb{E}[f(Z_i(t))]| = O(h)$$

and $\text{Var}(\widehat{Q}_Z(h)) = O(h)$ with $h < \varepsilon$. We will also need an approximate path to estimate the term $f(X_i(T)) - f(Z_i(T))$. This will be denoted by $Z_i^h(T)$ and taken to satisfy

$$\mathbb{E}[|Z_i^h(T) - Z_i(T)|^2] = O(h^2).$$

Note that h will generally depend on ε, and hence Ω, so the use of the $O(\cdot)$ notation in this context is consistent with our earlier statement that asymptotic results are with respect to the system size. Such a bound will be satisfied provided the numerical integration step in algorithm 2 is convergent. Then an $O(\varepsilon)$ estimator $\widehat{Q}_{\text{coupled}}(M_1, h)$ of the form

$$\mathbb{E}f(X_i(T)) = \mathbb{E}[f(X_i(T)) - f(Z_i^h(T))] + \mathbb{E}[f(Z_i^h(T))] \tag{4.2}$$

can be constructed by summing the estimator

$$\widehat{Q}_{(X,Z^h)}(M_1) = \frac{1}{M_1}\sum_{j=1}^{M_1}(f(X_{i,[j]}(T)) - f(Z_{i,[j]}^h(T)))$$

and the approximation $\widehat{Q}_Z(h)$:

$$\widehat{Q}_{\text{coupled}}(M_1, h) = \widehat{Q}_{(X,Z^h)}(M_1) + \widehat{Q}_Z(h).$$

Setting

$$V_{(X,Z_h)}(T) = \text{Var}(f(X_i(T)) - f(Z_i^h(T))), \tag{4.3}$$

the variance of the coupled MC estimator, V_{coupled}, is simply

$$V_{\text{coupled}} = M_1^{-1} V_{(X,Z_h)} + \text{Var}(\widehat{Q}_Z(h)) \sim M_1^{-1} V^{\Omega}_{(X,Z^h)} \tag{4.4}$$

where we are assuming the first term is leading order in Ω, meaning that the limiting factor in reducing the variance in the simulation of the coupled path. The methods of this paper are most applicable when $\widehat{Q}_Z(h)$ is computed with non-Monte Carlo based methods, and hence $\text{Var}(\widehat{Q}_Z(h)) = 0$. However, we do not exclude the possibility that $\widehat{Q}_Z(h)$ is computed from Monte Carlo simulations of the PDMP. In such cases we would assume that simulations of $Z^h(t)$ are much cheaper than simulations of $X(t)$ and hence $\text{Var}(\widehat{Q}_Z(h))$ can be made $o(\epsilon)$ at a negligible cost, for example, this will be the case when the search step in Algorithm 1 is very expensive. Of course for an order ε estimator, we require $V_{\text{coupled}} = O(\varepsilon^2)$ so that

$$M_1 = \varepsilon^{-2} V_{(X,Z_h)}.$$

It is now clear that for $\widehat{Q}_{\text{coupled}}$ to be preferable over $\widehat{Q}_{\text{crude}}$ we must be able to make $V_{(X,Z_h)}$ small and $\widehat{Q}_Z(h)$ must be cheap to generate. To be more precise, set

$$\begin{aligned} \mathsf{C}_{(X,Z^h)} &\approx \mathbb{E}[\ \#\text{ of computations to simulate the coupled process } (X, Z^h)\ ] \\ \mathsf{C}_{Q_{Z,h}} &\approx \mathbb{E}[\ \#\text{ of computations to compute } \widehat{Q}_Z(h)\ ] \end{aligned}$$

then the computational cost of the coupled MC estimator is

$$\mathsf{C}_{\text{coupled}} = M_1 \mathsf{C}_{(X,Z^h)} + \mathsf{C}_{\widehat{Q}_Z(h)}.$$

Generally we would expect $\mathsf{C}_{(X,Z^h)}$ to be comparable to C_X. One the other hand, while $\mathsf{C}_{\widehat{Q}_Z(h)}$ depends heavily on the method we use to generate $\widehat{Q}_Z(h)$, it will generally be true that $\mathsf{C}_{\widehat{Q}_Z(h)} \ll \mathsf{C}_X$ because generating statistics of $Z^h(t)$ does not require implementing the large number of stochastic events needed for $X(t)$. Moreover, as we have already mentioned it is often possible to bypass Monte Carlo methods entirely to obtain these statistics since the process $(Z(t), Y(t))$ my have a very low dimensional stochastic component. When Monte Carlo simulations are bypassed we only need to compute $\widehat{Q}_Z(h)$, even if we are interested in doing computational experiments over a range of values of Ω.

5. Models without feedback

For networks without feedback, we can simply couple the full stochastic model to the PDMP by generating a realization of $Y(t)$ and using it to drive both the full model and the PDMP. This simplifies the complexity analysis and also provides an upper bound on the achievable complexity of the coupled MCE for more general systems. Explicitly, the coupling is

$$Y(t) = Y(0) + \sum_{j \in \mathcal{J}_1} \Pi_j \left(\int_0^t \alpha_j(Y(s)) ds \right) u_j \tag{5.1a}$$

$$X(t) = X(0) + \sum_{j \in \mathcal{J}_0} \Pi_j \left(\Omega \int_0^t \alpha_j(X(s)/\Omega, Y(s)) ds \right) k_j \tag{5.1b}$$

$$Z(t) = Z(0) + \sum_{j \in \mathcal{J}_0} \int_0^t \alpha_j(Z(s), Y(s)) ds k_j \tag{5.1c}$$

It is not difficult to see that (5.1) is a probabilistic coupling in the sense that the marginal distributions of (X, Z) are indeed the distributions for the uncoupled processes $X(t)$ and $Z(t)$. In the present setting algorithm 2 can be replaced by the simpler algorithm 3 which reuses the process $Y(t)$ in an more efficient manner. Algorithm 3 is also easier to implement with existing software since once can use preexisting packages to perform algorithm 1 and the numerical integration.

The following Theorem, which is proved in Appendix B, gives a rigorous bound on the second moment of the pathwise error, which can be used to estimate the variance.

Algorithm 3 Simulation of (5.1)

1: Generate $\{Y(t)\}_{0\le t\le T}$ using algorithm 1.
2: Generate $\{X(t)\}_{0\le t\le T}$ using algorithm 1 with rates computed from $\{Y(t)\}_{0\le t\le T}$.
3: Generate $\{Z^h(t)\}_{0\le t\le T}$ using deterministic numerical integration with rates computed from $\{Y(t)\}_{0\le t\le T}$.

Theorem 5.1. *Suppose Assumptions 1 and 2 hold in a network without feedback and let* $(Y(t), X^\Omega(t), Z(t))$ *be given by (5.1). Then*

$$\mathbb{E}[||X^\Omega(T) - Z(T)||_1^2 \mathbf{1}\{\tau_K > T\}] = O(\Omega^{-1}). \tag{5.2}$$

Theorem 5.1 implies an upper bound on the variance of the coupled paths. To see this, note that Assumption 2 implies there exists $\bar{C}_2(T)$ indepndent of K and Ω such that.

$$\mathbb{E}[||X^\Omega(T) - Z(T)||_1^2 \mathbf{1}\{\tau_K < T\}] \le \bar{C}_2(T) C(T) K^{-1}.$$

Following an argument in [5], for each $\gamma \in (0,1)$ we can select K such that

$$\mathbb{E}[||X^\Omega(T) - Z(T)||_1^2 \mathbf{1}\{\tau_K < T\}] \le \gamma \mathbb{E}[||X^\Omega(T) - Z(T)||_1^2].$$

It follows that

$$\mathbb{E}[||X^\Omega(T) - Z(T)||_1^2] \le \frac{1}{1-\gamma} O(\Omega^{-1}) = O(\Omega^{-1}).$$

Letting L_f denote the Lipschitz constant of f, we have

$$\begin{aligned} V^\Omega_{(X,Z^h)} &= \mathrm{Var}(f(X_i^\Omega(T)) - f(Z_i^h(T))) \\ &\le 2\mathbb{E}[|f(X_i^\Omega(T)) - f(Z_i(T))|^2] + 2\mathbb{E}[|f(Z_i^h(T)) - f(Z_i(t))|^2] \\ &\le 2L_f^2 \mathbb{E}[||X^\Omega(T) - Z(T)||^2] + 2L_f^2 \mathbb{E}[||Z^h(T) - Z(T)||^2]. \end{aligned}$$

The first term is $O(h)$ (which is $o(\Omega^{-1})$) assuming our numerical integration scheme is convergent.

5.1. Complexity analysis

When assumptions 1 and 2 are satisfied in a network without feedback, equation (4.4) and Theorem 5.1 imply that

$$V_{\mathrm{cmc}} \sim M_1^{-1} V^\Omega_{(X,Z^h)} = M_1^{-1}\Omega^{-1}.$$

To make this asymptotically order ε^2 we set

$$M_1 = \Omega^{-1}\varepsilon^{-2} = \varepsilon^{1/\delta - 2}$$

The complexity of this estimator is then

$$\mathsf{C}_{\mathrm{coupled}} = \varepsilon^{1/\delta-2}\mathsf{C}_{(X,Z^h)} + \mathsf{C}_{\widehat{Q}_Z(h)}.$$

As noted earlier, the rare and abundant species can be simulated separately using algorithm 3. Defining

$$\begin{aligned} \mathsf{C}_Y &:= \text{Complexity of Step 1 in algorithm 3} \\ \mathsf{C}_{X|Y} &:= \text{Complexity of Step 2 in algorithm 3} \\ \mathsf{C}_{Z^h|Y} &:= \text{Complexity of Step 3 in algorithm 3} \end{aligned}$$

we have

$$\mathsf{C}_{(X,Z^h)} = \mathsf{C}_{X|Y} + \mathsf{C}_{Z^h|Y} + \mathsf{C}_Y.$$

On the other hand, for the crude MC estimator each path is obtained from the first two steps in algorithm 3, hence

$$\mathsf{C}_{\text{crude}} = \varepsilon^{-2}(\mathsf{C}_{X|Y} + \mathsf{C}_Y).$$

Therefore

$$\mathsf{C}_{\text{coupled}} = \varepsilon^{1/\delta}\mathsf{C}_{\text{crude}} + \varepsilon^{1/\delta-2}\mathsf{C}_{(Z^h|Y)} + \mathsf{C}_{\widehat{Q}_Z(h)}$$

Using a high order numerical integration scheme (for example, a Runga-Kutta method) renders $\mathsf{C}_{Z^h|Y}$ negligible at leading order, while $\mathsf{C}_{X|Y} = O(\varepsilon^{-1/\delta})$. Since the rates affecting $Y(t)$ do not scale with Ω, we have $\mathsf{C}_Y = O(1)$. This means that

$$\mathsf{C}_{\text{crude}} = O(\varepsilon^{-2-1/\delta})$$

while

$$\mathsf{C}_{\text{coupled}} = O(\varepsilon^{-2}) + \mathsf{C}_{Q_{Z,h}}.$$

We work under the realistic assumption that $\mathsf{C}_{Q_{Z,h}} = o(\varepsilon^{-2})$, and therefore $\mathsf{C}_{\text{coupled}} = O(\varepsilon^{-2})$.

6. Extension to models with feedback

Constructing an effective coupling for a model with feedback is more subtle. This is because we cannot simply reuse the realization $\{Y(t)\}_{0\le t\le T}$, as this would generally not yield a probabilistic coupling. In other words, if we let $(X^\Omega(t), Y^X(t))$ and $(Z(t), Y^Z(t))$ be given by (3.8) and (3.9) where we have replaced $Y(t)$ with $Y^X(t)$ and $Y^Z(t)$ to emphasize these are different paths, it is generally not the case that $Y^X(t)$ and $Y^Z(t)$ are equivalent in any standard sense of equivalence for stochastic processes. Formulating the problem in this way suggests that the appropriate generalization of the coupling (5.1) in this setting involves a coupling of $Y^X(t)$ and $Y^Z(t)$. Indeed, if we can construct a coupling that keeps $\Delta Y^\Omega(t) := Y^X(t) - Y^Z(t)$ small, then

$$(X^\Omega(t), Y^X(t)) = (X^\Omega(t), \Delta Y^\Omega(t) + Y^Z(t))$$

should remain close to $(Z(t), Y^Z(t))$, and when $\Delta Y^\Omega(t) = 0$ for all $t \ge 0$ we retrieve a coupling of the form (5.1).

The method we will use to couple $Y^X(t)$ and $Y^Z(t)$ is known in the probability literature as the *split coupling* [32] and has appeared in the context of coupling τ-leaping approximations with different values of τ [33, 32]. Let us introduce the split coupling in the context of a much simpler problem, namely the problem of coupling two homogenous Poisson processes with propensities α_1 and α_2. We continue to use the convention that $\Pi_i(t)$ are independent unit rate Poisson process, so the two processes we wish to couple can be written

$$Y_1(t) = \Pi_1(\alpha_1 t), \quad Y_2(t) = \Pi_2(\alpha_2 t).$$

The idea of the split coupling is to break up these reaction channels into two parts, one that is common to both the processes, and another that is specific to a given process. In this way, the coupling of $Y_1(t)$ and $Y_2(t)$ will be driven by three Poisson processes: one counting the events that occur in both $Y_1(t)$ and $Y_2(t)$, one counting the events occurring only in $Y_1(t)$, and another counting the events occurring only in $Y_2(t)$. In order to write down the explicit representations of these counting processes we introduce the function

$$\rho(\alpha_1, \alpha_2) := \alpha_1 - \alpha_2 \wedge \alpha_1.$$

Now we observe that for two Poisson processes $\Pi_3(t)$ and $\Pi_4(t)$

$$\Pi_3((\alpha_1 \wedge \alpha_2)t) + \Pi_4(\rho(\alpha_1, \alpha_2)t)$$

is identical in distribution to $Y_1(t)$ due to the additive property of the Poisson processes. It should then be clear that (after introducing one additional Poisson processe $\Pi_5(\cdot)$), a coupled representation $(Y_1(t), Y_2(t))$ of $Y_1(t)$ and $Y_2(t)$ is given by

$$\begin{aligned} Y_1(t) &= \Pi_3((\alpha_1 \wedge \alpha_2)t) + \Pi_4(\rho(\alpha_1, \alpha_2)t) \\ Y_2(t) &= \Pi_3((\alpha_1 \wedge \alpha_2)t) + \Pi_5(\rho(\alpha_2, \alpha_1)t) \end{aligned}$$

Here $\Pi_3(\cdot)$ is the common counting process described above that counts the events occurring in both $Y_1(t)$ and $Y_2(t)$, while $\Pi_4(\cdot)$ and $\Pi_5(\cdot)$ count the events only occurring in one process. If $\alpha_1 \approx \alpha_2$, then $\rho(\alpha_1, \alpha_2)$ and $\rho(\alpha_2, \alpha_1)$ will be small, so that there are relatively few events uncommon to both process, and hence $Y_1(t)$ and $Y_2(t)$ remain close.

Returning to the problem of coupling $Y_X(t)$ and $Y_Z(t)$, the split coupling can be constructed in exactly the same manner for each reaction channel indexed by $\mathcal{J}_1$. That is, for each $j \in \mathcal{J}_1$ we replace two uncoupled counting processes of the form

$$\Pi_{1,j}\left(\int_0^t \alpha_j(X^\Omega(s), Y^X(s))ds\right), \quad \Pi_{2,j}\left(\int_0^t \alpha_j(Z(s), Y^Z(s))ds\right),$$

where $\Pi_{1,j}(t)$ and $\Pi_{2,j}(t)$ are independent unit rate Poisson process, with

$$\begin{aligned}
&\Pi_{3,j}\left(\int_0^t \alpha_j(Z(s), Y^Z(s)) \wedge \alpha_j(X^\Omega(s), Y^X(s))ds\right) \\
&\qquad + \Pi_{4,j}\left(\Omega \int_0^t \rho(\alpha_j(X^\Omega(t), Y^X(s)), \alpha_j(Z(s), Y^Z(s)))ds\right) \\
&\Pi_{3,j}\left(\int_0^t \alpha_j(Z(s), Y^Z(s)) \wedge \alpha_j(X^\Omega(s), Y^X(s))ds\right) \\
&\qquad + \Pi_{5,j}\left(\Omega \int_0^t \rho(\alpha_j(Z(s), Y^Z(s)), \alpha_j(X^\Omega(t), Y^X(s)))ds\right)
\end{aligned}$$

where $\Pi_{3,j}(t)$, $\Pi_{4,j}(t)$ and $\Pi_{5,j}(t)$ are all independent unit rate Poisson processes.

After substituting the coupled reaction channels into (3.8b) and (3.9b), with $\Pi_j \to \Pi_{1,j}$ and $\Pi_j \to \Pi_{2,j}$, respectively, we obtain their coupled representation

$$\begin{aligned}
Y^X(t) = Y(0) &+ \sum_{j\in\mathcal{J}_1} \Pi_{3,j}\left(\int_0^t \alpha_j(Z(s), Y^Z(s)) \wedge \alpha_j(X^\Omega(t), Y^X(s))ds\right) u_j \\
&+ \sum_{j\in\mathcal{J}_1} \Pi_{4,j}\left(\int_0^t \rho(\alpha_j(X^\Omega(s), Y^X(s)), \alpha_j(Z(s), Y^Z(s)))ds\right) u_j
\end{aligned} \tag{6.1a}$$

$$\begin{aligned}
Y^Z(t) = Y(0) &+ \sum_{j\in\mathcal{J}_1} \Pi_{3,j}\left(\int_0^t \alpha_j(Z(s), Y^Z(s)) \wedge \alpha_j(X^\Omega(t), Y^X(s))ds\right) u_j \\
&+ \sum_{j\in\mathcal{J}_1} \Pi_{5,j}\left(\int_0^t \rho(\alpha_j(Z(s), Y^Z(s)), \alpha_j(X^\Omega(s), Y^X(s)))ds\right) u_j
\end{aligned} \tag{6.1b}$$

$$X^\Omega(t) = X^\Omega(0) + \sum_{j\in\mathcal{J}_0} \Pi_j\left(\Omega \int_0^t \alpha_j(X^\Omega(s), Y^X(s))ds\right) k_j \tag{6.1c}$$

$$Z(t) = Z(0) + \sum_{j\in\mathcal{J}_0} \int_0^t \alpha_j(Z(s), Y^Z(s))ds k_j. \tag{6.1d}$$

The structure of the coupled process is similar to that of the PDMP (3.9) in that there is a piecewise deterministic component, $Z(t)$, driven by a jump process whose evolution is governed by (6.1a-6.1c). In other words, (6.1) can be rewritten in the form of (3.9) by replacing $Y(t)$ with $(Y^X(t), Y^Z(t), X(t))$ and adjusting the rates accordingly. Viewing the coupled paths in this way is convenient since we already know how to simulate a generic PDMP. In particular, to generate sample paths of the coupled process (6.1) we simply use Algorithm 2 with the appropriate replacements.

In order to extend Theorem 5.1 to the coupled processes in equation (6.1), recall that our original motivation for coupling $Y^X(t)$ and $Y^Z(t)$ was so that the resulting coupled process would look like a

perturbation of the coupling (5.1). Specifically, the perturbations come from the poisson processes $\Pi_{4,j}(\cdot)$ and $\Pi_{5,j}(\cdot)$, and if these terms vanish, $Y^X(t) = Y^Z(t)$, so that we retrieve the coupling of equation (5.1). Using the process $\Delta Y^\Omega(t)$ defined above, we can write equations (6.1c) and (6.1d) as

$$X^\Omega(t) = X^\Omega(0) + \sum_{j\in\mathcal{J}_0} \Pi_j\left(\Omega\int_0^t \alpha_j(X^\Omega(s), Y^X(s))ds\right)k_j$$
$$Z(t) = Z(0) + \sum_{j\in\mathcal{J}_0} \int_0^t \alpha_j(Z(s), Y^X(s) + \Delta Y^\Omega(s))k_j ds.$$

In the proof of Theorem 5.1 we relied on a Lipschitz bound of the form (3.3). Recalling the properties of the rates α_j, the appropriate generalization of the Lipschitz bound in the present context is

$$|\alpha_j(x_1, y) - \alpha_j(x_2, y + \Delta y)| \le L_{\alpha_j}(K)||x_1 - x_2|| + \sum_{p=1}^{\deg(\zeta_j(y))} a_{j,p}||x_2||^p||\Delta y||^p.$$

where $a_{j,p}$ are constants independent of Ω. Squaring this inequality,

$$\begin{aligned}&|\alpha_j(x_1, y) - \alpha_j(x_2, y + \Delta y)|^2\\ &\quad\le L_{\alpha_j}(K)^2||x_1 - x_2||^2 + 2\sum_{p=1}^{\infty} a_{j,p}L_{\alpha_j}(K)||x_1 - x_2||^p||x_2||^p||\Delta y||^p + \left(\sum_{p=1}^{\deg(\zeta_j(y))} a_{j,p}||x_2||^p||\Delta y||^p\right)^2\end{aligned} \tag{6.2}$$

It is clear that we now need to bound not only the moments of $||\Delta Y^\Omega(t)||_1$, as well as their products with $||X^\Omega(t) - Z(t)||_1$, in order to obtain a bound on the complexity of the coupled Monte Carlo estimator. Note that $Z(t)$ is bounded independent of Ω, so the $||x_2||$ terms in (6.2) do not play a significant role. For simplicity, we focus on the case where $\zeta_j(y)$ is linear ($\deg(\zeta_j(y)) = 1$), while the general case will be analyzed in a forthcoming more technical paper. The following Lemma, which is proved in Appendix C, provides us with the result we need.

Lemma 6.1. *Suppose Assumptions 1 and 2 hold and let $(Y^X(t), Y^Z(t), X^\Omega(t), Z(t))$ be given by (6.1). Assume $\zeta_j(y)$ is linear for each $j \in \mathcal{J}_2$. Then*

$$\mathbb{E}[||\Delta Y^\Omega(T)||_1^2\mathbf{1}\{\tau_K > T\}] = O(\Omega^{-1/2}) \tag{6.3a}$$
$$\mathbb{E}[||X^\Omega(T) - Z(t)||_1||\Delta Y^\Omega(T)||_1\mathbf{1}\{\tau_K > T\}] = O(\Omega^{-1/2}) \tag{6.3b}$$

Using Lemma 6.1 and (6.2), we can obtain the following result.

Theorem 6.1. *Suppose Assumptions 1 and 2 hold and let $(Y^X(t), Y^Z(t), X^\Omega(t), Z(t))$ be given by (6.1). Moreover, assume $\zeta_j(y)$ is linear for each $j \in \mathcal{J}_2$. Then*

$$\mathbb{E}[||X^\Omega(T) - Z(T)||_1^2\mathbf{1}\{\tau_K > T\}] = O(\Omega^{-1/2}). \tag{6.4}$$

Proof 6.1. *We get this result after inserting (6.2) into the proof of Theorem 6.4 and applying Lemma 6.1.*

It follows that $V^\Omega_{(X,Z^h)} = O(\Omega^{-1/2})$ when $X(t)$ and $Z^h(t)$ are coupled according to (6.1).

6.1. Complexity analysis

In light of the analysis for the coupled MCE without feedback, and Theorem (6.4), we would expect that $O(\varepsilon^{-2-1/2\delta})$ complexity is achieved in the presence of feedback. To obtain this result one only needs to show that $\mathsf{C}_{(X,Z^h)} = O(\Omega)$ where the path is obtained from Algorithm 2. It is safe to assume that the

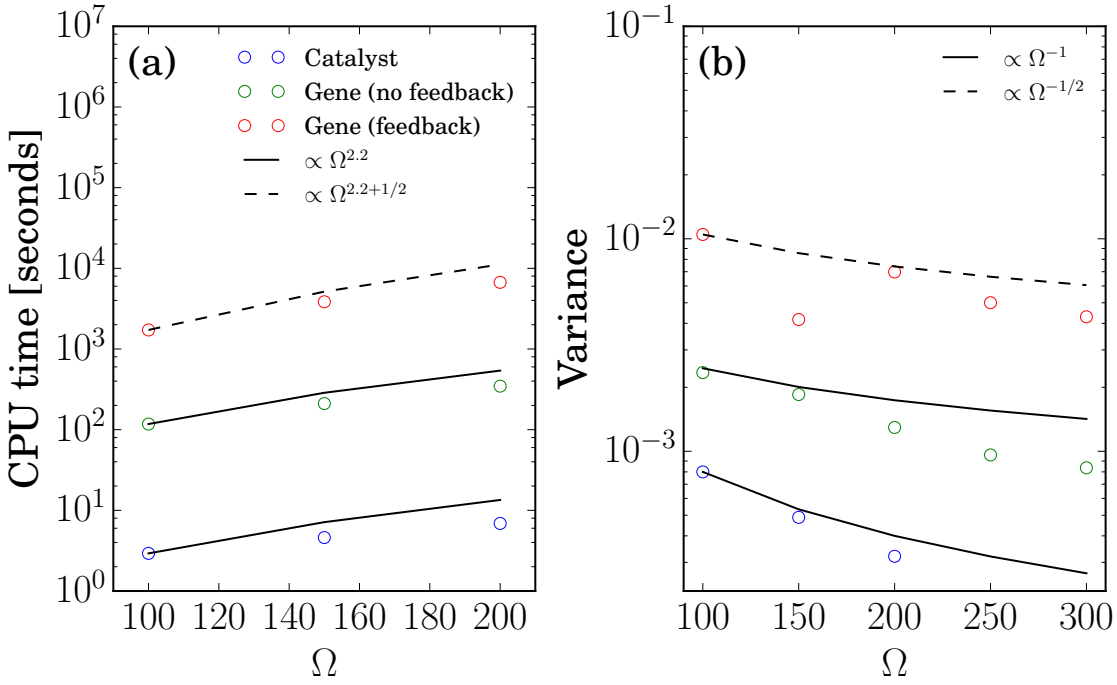

Figure 1: (a) The CPU time used to compute an estimate of $\mathbb{E}[X^{\Omega}(10)]$ for each example, as well as the theoretical bounds. We have used $\delta = 1.1$. (b) The estimated variances using 10^5 samples of the coupled paths, along with the theoretical bounds. Note the log scales on the time axis. We have used the LSODA algorithm to perform the numerical integration of (3.13). Parameter values used are $\alpha = 1$, $\beta = 1$, $\lambda = 1$. $\mu = 0.8$ for Example 3.2; and $\alpha = 1$, $\beta = 1$, $\lambda = 4$. $\mu = 5$. for Examples 3.3 and 3.4.

contribution to the complexity of the numerical integration steps in Algorithm 2 is $o(\Omega)$, but since we are solving for the coupled paths (6.1) the frequency of the jump events computed by solving 3.13 scales with Ω. These events then come to dominate the cost of this Algorithm so that we have $\mathsf{C}_{(X,Z^h)} = O(\Omega)$, and as expected, $C_{\text{coupled}} = O(\varepsilon^{-2-1/2\delta})$. Note that this result implies the computations are sped up by a factor of $O(\varepsilon^{-1/2\delta})$ in contrast to the $O(\varepsilon^{-1/\delta})$ speed up when there is no feedback.

7. Numerical Examples

To illustrate that our theoretical complexity results can be achieved in practice we have applied our coupled MCE to estimate $\mathbb{E}[X^{\Omega}(20)]$ for a range of system sizes. The results for Examples 3.2, 3.3 and 3.4 are displayed in Figure 1, where we have used the couplings (5.1) for the first two examples, and (6.1) for the third.[1] For Examples 3.2 we have solved for $\widehat{Q}_Z(h)$ by integrating the moment flow equations (A.4), while in Examples 3.2 and 3.4 we have used the solution to (A.2) given in [22].

8. Conclusions

Variance reduction in Monte Carlo estimators through probabilistic coupling has been used extensively in the scientific computing literature [33, 32, 15, 16]. However, there has been little work exploring the application of simplified models to reduce variances in Monte Carlo estimators for complex chemical reaction

[1] These computations were performed on a 2014 MacBook Pro with a 2.6 GHz Intel Core i5 processor using Python 3

networks. We have extended the idea of variance reduction to models with partial thermodynamic limits in which the qualitative behavior of the full stochastic model is well approximated by a PDMP. Such population models arise in the biological and chemical sciences whenever the population can be decomposed into a group of abundant species, and a group of rare species. The rare species often act as an environment that controls the dynamics of a large population, such as how the discrete state of a gene controls the production of a protein. When there is no feedback from the abundant species in the population, the coupling of the full model to the PDMP is straightforward. and we have shown that the coupling is indeed effective in reducing the variance of the Monte Carlo estimator. We have also derived an effective coupling for models with feedback, which applies to an extremely wide class of population models.

Our results suggest that approximate stochastic models, such as the ones studied rigorously in [4, 2, 5] may be useful in the context of variance reduction for exact models. It would be particularly fruitful to extend our work to develop computational tools that are specifically tailored to spatial process. In particular, the reaction diffusion master equation (RDME) is a continuous time Markov chain approximation of reaction diffusion processes for which there is a great deal of interest in simulating efficiently [34, 35]. Other future directions include extending the coupling to other model reductions, such as the quasi-steady state, an idea that was briefly explored in [15].

Acknowledgements

This work was supported by the National Science Foundation (DMS 1613048 and RTG 1148230)

A. Non-Monte Carlo based methods for a PDMP

The efficiency of the Monte Carlo methods introduced in this paper depend on our ability to obtain statistics of the PDMP without resorting to Monte-Carlo based methods. In this section we give a brief overview of the non-Monte Carlo based analysis of the PDMP. All non-Monte Carlo based methods are derived from the forward Kolmogorov equation. While the Kolmogorov equation for the full model may be prohibitively complex due to the combinatorial complexity when there are many particles, the PDMP is often very simple. We note that the problem of developing efficient numerical methods for solving the Kolmogorov equation for the full model and the PDMP is a very active area of research, but beyond the scope of this work.

The main object of interest will be the generator of the process $(Z(t), Y(t))$,

$$\begin{aligned} Af(z,y) = &\left(\sum_{j\in\mathcal{J}_0}\alpha_j(z,y)k_j\right)\cdot\nabla_z f(z,y) \\ &+\sum_{j\in\mathcal{J}_1}\alpha_j(y-u_j)f(z,y-u_j)-\alpha_j(y)f(z,y). \end{aligned} \tag{A.1}$$

A.1. Differential Chapman Kolmogorov equations

The forward Kolmogorov equation, or *differential Chapman Kolmogorov* (dCK) equation

$$\frac{d}{dt}p(z,y,t) = A^\dagger p(z,y,t) \tag{A.2}$$

gives the evolution of the density

$$p(z,y,t) = \mathbb{P}(Z(t)=z, Y(t)=y).$$

Here the superscript † denotes the Hilbert adjoint operator with respect to the appropriate inner product. The dCK equations are a system of hyperbolic PDEs, and will generally only be useful for numerical simulation when they are finite dimensional. This occurs exactly when the rare species are conserved ($||Y||$

is constant) and the methods presented in this paper apply when this condition holds. The situation is also tricky when the abundant species are not conserved because the domain of (A.2) may be unbounded. However, we can often solve the PDE on a bounded domain provided the deterministic dynamics have a trapping region. For example, see [22].

For example 3.2 with $c = 1$ the generator is

$$\begin{aligned} Af(z,0) &= -\alpha z \partial_z f(z,0) - \lambda f(z,0) + \mu f(z,1) \\ Af(z,1) &= \beta(1-z)\partial_z f(z,1) + \lambda f(z,0) - \mu f(z,1) \end{aligned}$$

so the dCK equation is

$$\begin{aligned} \frac{d}{dt}p(z,0,t) &= \alpha \partial_z[z p(z,0,t)] - \lambda p(z,0,t) + \mu p(z,1,t) \\ \frac{d}{dt}p(z,1,t) &= -\beta \partial_z[(1-z) p(z,1,t)] + \lambda p(z,0,t) - \mu p(z,1,t), \end{aligned}$$

along with the reflecting boundary conditions

$$p(0,y,t) = p(1,y,t) = 0.$$

A.2. Moment flow equations

While solving (A.2) numerically is preferable compared to Monte Carlo simulations of the PDMP (3.9) for chemical reaction networks, significant computational gains can be made when one is only interested in expected values of polynomial function of the process $Z(t)$. This is done by deriving moment equations which are ODEs for the dynamic evolution of the moments and point correlation of $(Z(t), Y(t))$. We derive these from the backward Kolmogorov (bK) equation

$$\frac{d}{dt}\mathbb{E}[f(Z(t),Y(t))] = \mathbb{E}[Af(Z(t),Y(t))]. \tag{A.3}$$

For notationally simplicity we derive the moment equations when $z \in \mathbb{R}$, although this derivation can easily be generalized [36]. It is not difficult to see that the polynomial (3.4) is invariant under A. In particular, if we set $f^{(r)}_{y'}(z,y) = z^r \mathbf{1}\{y = y'\}$ and write the polynoomial in the form

$$\alpha_j(x,y) = \zeta_j(y) \sum_{i=1}^{\deg(\alpha_j)} \alpha_{j,i} x^i,$$

then

$$\begin{aligned} Af^{(r)}_{y'}(z,y) &= \sum_{j \in \mathcal{J}_0} \alpha^r_j(z,y) k_j r f^{(r-1)}_{y'}(z,y) + \sum_{j \in \mathcal{J}_1} \alpha_j(y-u_j) f^{(r)}_{y'}(z,y-u_j) - \alpha_j(y) f^{(r)}_{y'}(z,y) \\ &= \sum_{j \in \mathcal{J}_0} \zeta_j(y) r k_j \sum_{i=1}^{\deg(\alpha_j)} \alpha_{j,i} f^{(r-1+i)}_{y'}(z,y) \\ &+ \sum_{j \in \mathcal{J}_1} \zeta_j(y-u_j) \sum_{i=1}^{\deg(\alpha_j)} \alpha_{j,i} f^{(r+i)}_{y'}(z,y-u_j) - \zeta_j(y) \sum_{i=1}^{\deg(\alpha_j)} \alpha_{j,i} f^{(r+i)}_{y'}(z,y) \end{aligned}$$

This expression for $Af^{(r)}_{y'}(z,y)$ is particularly useful when A preserves the degree of polynomials. This will happen when there is no feedback and $\deg(\alpha_j) = 1$. Then, setting $m^{(r)}_y(t) = \mathbb{E}[f^{(r)}_y(Z(t),Y(t))]$, taking the

exception of $Af_{y'}^{(r)}(Z(t),y)$ with respect to Z, and applying the bK equation (A.3) yields the closed system of linear differential equations,

$$\begin{aligned}\frac{d}{dt}m_y^{(r)}(t) = &\sum_{j\in\mathcal{J}_0}\zeta_j(y)rk_j\sum_{i\in\{0,1\}}\alpha_{j,i}m_y^{(r-1+i)}(t)\\ &+\sum_{j\in\mathcal{J}_1}\zeta_j(y-u_j)m_{y+u_i}^{(r)}(t)-\zeta_j(y)m_y^{(r)}(t)\end{aligned} \tag{A.4}$$

which we refer to as the *moment flow equations*. If $\deg(\alpha_j)>1$ or the reaction network has feedback, then the system is not closed and one has to resort to moment closure techniques, which generally approximate the infinite-dimensional system of linear equations by a finite-dimensional system of nonlinear equations. See for example [36, 37] for one approach. The efficiency of implementing a moment closure scheme depends heavily on the specific problem, and a detailed discussion of the computational issues associated with solving coupling moment equations is beyond the scope of this work.

One complication that arises in deriving moment equations is that the conservation rules used to derive the dCK equation may be applicable to the moment equations. For example, taking $c=1$ in example 3.2 we need to derive the moment flow equations for the first moments

$$m_{i,j}^{(1)}(t)=\mathbb{E}[X_i(t)\mathbf{1}Y(t)=j]$$

where $X_i(t)=i+X(t)(1-2i)$, $i\in\{0,1\}$. These are

$$\begin{aligned}\frac{d}{dt}m_{0,0}^{(1)}(t)&=-\alpha m_{0,0}^{(1)}(t)+\lambda m_{0,1}^{(1)}(t)-\mu m_{0,0}^{(1)}(t)\\ \frac{d}{dt}m_{0,1}^{(1)}(t)&=\beta m_{1,1}^{(1)}(t)-\lambda m_{0,1}^{(1)}(t)+\mu m_{0,0}^{(1)}(t)\\ \frac{d}{dt}m_{1,0}^{(1)}(t)&=\alpha m_{0,0}^{(1)}(t)+\lambda m_{1,1}^{(1)}(t)-\mu m_{1,0}^{(1)}(t)\\ \frac{d}{dt}m_{1,1}^{(1)}(t)&=-\beta m_{1,1}^{(1)}(t)-\lambda m_{1,1}^{(1)}(t)+\mu m_{1,0}^{(1)}(t).\end{aligned}$$

The original system has two conservation rules but the moment equations only have one, so the system above could be reduced to a three dimensional system, but we cannot obtain a two dimensional system of moment equations from the two dimensional dCK equation. Instead we have derived these from the dCK equations without conservation.

B. Proof of Theorem 5.2

Throughout this Appendix $||\cdot||=||\cdot||_1$ and $t<T$ is fixed. Let $L_{\alpha_j}(K)$ denote the maximum Lipschitz constant of $\alpha_j(\cdot,y)$ over all values of y, which is independent of Ω but depends quadratically on K. Define the constants

$$\begin{aligned}B_1&=\sum_{j\in\mathcal{J}_0}||k_j||L_{\alpha_j}(K)\\ B_2&=\max_j||k_j||\mathcal{J}_0|\end{aligned}$$

and let $\widehat{\Pi}_i:=\Pi_i(t)-t$ denote the centered Poisson process. Note that both B_1 and B_2 are independent of Ω.

Proof B.1 (Proof of Theorem 5.2). *Using the triangle inequality and adding and subtracting the appropriate terms, we have*

$$\begin{aligned}
||X^{\Omega}(t) - Z(t)|| &\leq \max_{j\in\mathcal{J}_0}||k_j|| \sum_{j\in\mathcal{J}_0} \left|\Omega^{-1}\Pi_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right) - \int_0^t \alpha_j(Z(s),Y(s))ds\right| \\
&= \max_{j\in\mathcal{J}_0}||k_j|| \sum_{j\in\mathcal{J}_0} \Bigg|\int_0^t \alpha_j(X^{\Omega}(s),Y(s)) - \alpha_j(Z(s),Y(s))ds \\
&\qquad + \Omega^{-1}\widehat{\Pi}_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\Bigg| \\
&\leq \max_{j\in\mathcal{J}_0}||k_j|| \sum_{j\in\mathcal{J}_0} \int_0^t |\alpha_j(X^{\Omega}(s),Y(s)) - \alpha_j(Z(s),Y(s))|ds \\
&\qquad + \max_{j\in\mathcal{J}_0}||k_j|| \sum_{j\in\mathcal{J}_0} \Omega^{-1}\left|\widehat{\Pi}_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\right| \\
&\leq \max_{j\in\mathcal{J}_0}||k_j|| \sum_{j\in\mathcal{J}_0} \int_0^t L_{\alpha_j}(K)||X^{\Omega}(s) - Z(s)||ds \\
&\qquad + \max_{j\in\mathcal{J}_0}||k_j|| \sum_{j\in\mathcal{J}_0} \Omega^{-1}\left|\widehat{\Pi}_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\right| \\
&\leq B_1\int_0^t ||X^{\Omega}(s) - Z(s)||ds + \max_{j\in\mathcal{J}_0}||k_j|| \sum_{j\in\mathcal{J}_0} \Omega^{-1}\left|\widehat{\Pi}_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\right|.
\end{aligned}$$

Then squaring both sides of our inequality and applying Cauchy-Schwarz yields

$$\begin{aligned}
||X^{\Omega}(t) - Z(t)||^2 \leq 3B_1^2 t\int_0^t ||X^{\Omega}(s) - Z(s)||^2 ds \\
+ 3B_2^2\Omega^{-2}\sum_{j\in\mathcal{J}_0}\left(\widehat{\Pi}_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\right)^2.
\end{aligned}$$

Since the quadratic variation of the Martingale $\widehat{\Pi}_i(\cdot)$ is $\Pi_j(\cdot)$ and $\widehat{\Pi}_i(\cdot)^2 - \Pi_j(\cdot)$ is Martingale [19], the optional stopping theorem implies

$$\mathbb{E}\left[\left(\widehat{\Pi}_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\right)^2\right] = \mathbb{E}\left[\Pi_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\right],$$

and noting that $t \leq \tau_K$ an application of the smoothing formula [31] gives

$$\mathbb{E}\left[\Pi_j\left(\Omega\int_0^t \alpha_j(X^{\Omega}(s),Y(s))ds\right)\right] \leq \mathbb{E}\left[\Pi_j\left(\Omega t L_{\alpha_j}(K)K\right)\right] = \Omega t L_{\alpha_j}(K)K$$

where K is given in (3.6). Now set

$$B_3 = 4T^2\sum_{j\in\mathcal{J}_0} L_{\alpha_j}(K).$$

Taking the expectation of the original expression and applying Fubini's theorem we get

$$\mathbb{E}[||X^{\Omega}(t) - Z(t)||^2] \leq B_1\int_0^t \mathbb{E}[||X^{\Omega}(s) - Z(s)||^2]ds + \Omega^{-1}B_2^2 B_3 K.$$

After noting that all constants are $O(1)$, the result follows from Grönwall's inequality.

C. Proof of Lemma 6.3

In this Appendix $a \leq_C b$ will indicate that $a \leq Cb$ for some constant C that is independent of Ω (but may depend on K and t). Note that we are referring to different constants C in each use of this notation. In order to increase readability we will not keep track of the specific values of C, but the proof of Theorem 5.2 gives a blueprint for how these constants can be derived. In the following proofs we will refer to the random variables

$$A_{1,j}(t) := \int_0^t \rho(\alpha_j(X^\Omega(s), Y^X(s)), \alpha_j(Z(s), Y^Z(s)))ds$$
$$A_{2,j}(t) := \int_0^t \rho(\alpha_j(Z(s), Y^Z(s)), \alpha_j(X^\Omega(s), Y^X(s)))ds.$$

Of course

$$\Delta Y^\Omega(t) = \sum_{j \in \mathcal{J}_2} \left(\Pi_{4,j}(A_{1,j}(t)) - \Pi_{5,j}(A_{2,j}(t))\right) u_j.$$

Lemma C.1. *Suppose Assumptions 1 and 2 hold and let $(Y^X(t), Y^Z(t), X^\Omega(t), Z(t))$ be given by (6.1). Moreover, assume $\zeta_j(y)$ is linear for each $j \in \mathcal{J}_2$. Then*

$$\mathbb{E}[||X^\Omega(t) - Z(t)||\mathbf{1}\{\tau_K > T\}] \leq_C \int_0^t \mathbb{E}[||\Delta Y^\Omega(s)||]ds + O(\Omega^{-1/2})$$

Proof C.1. *Take $t < \tau_K$. Following the proof of Theorem 5.2, we have*

$$||X^\Omega(t) - Z(t)|| \leq_C \int_0^t ||X^\Omega(s) - Z(s)||ds + \int_0^t ||\Delta Y^\Omega(s)||||Z(s)||ds + \sum_{j \in \mathcal{J}_0} \Omega^{-1} \left|\widehat{\Pi}_j\left(\Omega \int_0^t \alpha_j(X^\Omega(s), Y(s))ds\right)\right|.$$

Using Jenson's inequality and the Martingale stopping theorem,

$$\begin{aligned}\mathbb{E}\left[\left|\widehat{\Pi}_j\left(\Omega \int_0^t \alpha_j(X^\Omega(s), Y(s))ds\right)\right|\right] &\leq \mathbb{E}\left[\left|\widehat{\Pi}_j\left(\Omega t L_{\alpha_j} K t\right)\right|\right] \leq \left(\mathbb{E}\left[\left(\widehat{\Pi}_j\left(\Omega t L_{\alpha_j} K t\right)\right)^2\right]\right)^{1/2} \\ &= \left(\mathbb{E}\left[\Pi_j\left(\Omega t L_{\alpha_j} K t\right)\right]\right)^{1/2} = O(\Omega^{1/2})\end{aligned}$$

from which is follows that

$$\mathbb{E}[||X^\Omega(t) - Z(t)||] \leq_C \int_0^t \mathbb{E}[||X^\Omega(s) - Z(s)||]ds + \int_0^t \mathbb{E}[||\Delta Y^\Omega(s)||||Z(s)||]ds + O(\Omega^{-1/2})$$

so that the result is obtained from Grönwall's inequality and the fact that $Z(t)$ is bounded and independent of Ω.

Lemma C.2. *Suppose Assumptions 1 and 2 hold and let $(Y^X(t), Y^Z(t), X^\Omega(t), Z(t))$ be given by (6.1). Moreover, assume $\zeta_j(y)$ is linear for each $j \in \mathcal{J}_2$. Then*

$$\mathbb{E}[||\Delta Y^\Omega(t)||\mathbf{1}\{\tau_K > T\}] \leq_C \int_0^t \mathbb{E}[||X^\Omega(s) - Z(s)||]ds.$$

Proof C.2. *Take $t < \tau_K$.*

$$||\Delta Y^{\Omega}(t)||_1 \leq_C \sum_{j \in \mathcal{J}_1} |\Pi_{4,j}\left(A_{1,j}(t)\right) - \Pi_{5,j}\left(A_{2,j}(t)\right)|$$

Since the Poisson processes are positive,

$$\begin{aligned}\mathbb{E}[|\Pi_{4,j}\left(A_{1,j}(t)\right) - \Pi_{5,j}\left(A_{2,j}(t)\right)|] &\leq \mathbb{E}[\Pi_{4,j}\left(A_{1,j}(t)\right) + \Pi_{5,j}\left(A_{2,j}(t)\right)] \\ &= \mathbb{E}[A_{1,j}(t) + A_{2,j}(t)].\end{aligned}$$

Now observe that for any $\alpha_1, \alpha_2 > 0$,

$$\begin{aligned}\rho(\alpha_1, \alpha_2) + \rho(\alpha_2, \alpha_1) &= \alpha_1 - \alpha_1 \wedge \alpha_2 + \alpha_2 - \alpha_1 \wedge \alpha_2 \\ &= \alpha_1 + \alpha_2 - 2\alpha_1 \wedge \alpha_2 \\ &\leq 2(\alpha_1 \vee \alpha_2 - \alpha_1 \wedge \alpha_2) = 2|\alpha_1 - \alpha_2|,\end{aligned}$$

and hence

$$\begin{aligned}A_{1,j}(t) + A_{2,j}(t) &\leq \int_0^t \rho(\alpha_j(X^{\Omega}(s), Y^X(s)), \alpha_j(Z(s), Y^Z(s))) \\ &\quad + \rho(\alpha_j(Z(s), Y^Z(s)), \alpha_j(X^{\Omega}(s), Y^X(s)))ds \\ &\leq 2\int_0^t |\alpha_j(X^{\Omega}(s), Y^X(s)) - \alpha_j(Z(s), Y^Z(s)))| \\ &\leq_C \int_0^t ||X^{\Omega}(s) - Z(s)||ds + \int_0^t ||\Delta Y^{\Omega}(s)||||Z(s)||ds.\end{aligned}$$

Consolidating our findings,

$$\mathbb{E}[||\Delta Y^{\Omega}(t)||] \leq_C \int_0^t \mathbb{E}[||X^{\Omega}(s) - Z(s)||] + \int_0^t \mathbb{E}[||\Delta Y^{\Omega}(s)||||Z(s)||]ds.$$

Grönwall's inequality once again yields the final result.

Theorem C.1. *Suppose Assumptions 1 and 2 hold in a network with feedback and let $(Y^X(t), Y^Z(t), X^{\Omega}(t), Z(t))$ be given by (6.1). Moreover, assume $\zeta_j(y)$ is linear for each $j \in \mathcal{J}_2$.*

$$\mathbb{E}[||\Delta Y^{\Omega}||_1 \mathbf{1}\{\tau_K > T\}] = O(\Omega^{-1/2}) \tag{C.1a}$$

$$\mathbb{E}[||X^{\Omega}(T) - Z(T)||_1 \mathbf{1}\{\tau_K > T\}] = O(\Omega^{-1/2}) \tag{C.1b}$$

Proof C.3. *Combing Lemma C.1 and Lemma C.2, we have*

$$\mathbb{E}[||\Delta Y^{\Omega}(t)||_1] \leq_C \int_0^t \mathbb{E}[||\Delta Y^{\Omega}(s)||_1]ds + O(\Omega^{-1/2})$$

so that Grönwall's inequality yields (C.1a) and (C.1b) follows from another application of Lemma C.2.

Lemma C.3. *Let $\Pi_1(t)$ and $\Pi_2(t)$ be $\mathcal{F}_t$-adapted unit rate Poisson processes and t_1 and t_2 be $\mathcal{F}_t$-stopping times with $t_1 \vee t_2 \leq D$ almost surely. Then,*

$$\mathbb{E}[\Pi_1(t_1)^2 + \Pi_2(t_2)^2] \leq (2D + 1)\mathbb{E}[t_1 + t_2].$$

Proof C.4. *From [5, Lemma 3.1] we have*

$$\mathbb{E}[\Pi_j(t_j)^2] = 2\mathbb{E}[\Pi_j(t_j)t_j] - \mathbb{E}[t_j^2] + \mathbb{E}[t_j],$$

from which it follows that

$$\begin{aligned}
&\mathbb{E}[\Pi_1(t_1)^2 + \Pi_2(t_2)^2] \\
&\quad = 2\mathbb{E}[\Pi_1(t_1)t_1 + \Pi_2(t_2)t_2] - \mathbb{E}[t_1^2 + t_2^2] + \mathbb{E}[t_1 + t_2] \\
&\quad \leq 2\mathbb{E}[(t_1 \vee t_2)(\Pi_1(t_1) + \Pi_2(t_2))] + \mathbb{E}[t_1 + t_2] \\
&\quad \leq 2D\mathbb{E}[\Pi_1(t_1) + \Pi_2(t_2)] + \mathbb{E}[t_1 + t_2] = (2D+1)\mathbb{E}[t_1 + t_2].
\end{aligned}$$

Proof C.5 (Proof of Lemma 6.1). *We first establish (6.3a):*

$$\begin{aligned}
||\Delta Y^\Omega(t)||^2 &\leq_C \sum_{j\in\mathcal{J}_2} \left(\Pi_{4,j}\left(A_{1,j}(t)\right) - \Pi_{5,j}\left(A_{2,j}(t)\right)\right)^2 \\
&\leq_C \sum_{j\in\mathcal{J}_2} \left(\Pi_{4,j}\left(A_{1,j}(t)\right) + \Pi_{5,j}\left(A_{2,j}(t)\right)\right)^2 \\
&\leq_C \sum_{j\in\mathcal{J}_2} \Pi_{4,j}\left(A_{1,j}(t)\right)^2 + \Pi_{5,j}\left(A_{2,j}(t)\right)^2
\end{aligned}$$

In order to apply Lemma C.3 we need to bound the internal times. For $A_{1,j}(t)$,

$$\begin{aligned}
A_{1,j}(t) &= \int_0^t \alpha_j(X^\Omega(s), Y^X(s)) - \alpha_j(X^\Omega(s), Y^X(s)) \wedge \alpha_j(Z(s), Y^Z(s)))ds \\
&\leq \int_0^t \alpha_j(X^\Omega(s), Y^X(s))ds \leq t\alpha_j(K, c).
\end{aligned}$$

In fact, the same bound holds for $A_{2,j}(t)$. It follows that

$$\begin{aligned}
\mathbb{E}[\Pi_{4,j}\left(A_{1,j}(t)\right)^2 + \Pi_{5,j}\left(A_{2,j}(t)\right)^2] &\leq (t\alpha_j(K,c)+1)\mathbb{E}[\Pi_{4,j}\left(A_{1,j}(t)\right) + \Pi_{5,j}\left(A_{2,j}(t)\right)] \\
&= (t\alpha_j(K,c)+1)\mathbb{E}[A_{1,j}(t) + A_{2,j}(t)] \\
&\leq_C \int_0^t \mathbb{E}[||X^\Omega(s) - Z(s)||]ds + \int_0^t \mathbb{E}[||\Delta Y^\Omega(s)||]ds
\end{aligned}$$

Theorem C.1 now implies (6.3a). In order to obtain (6.3b), simply note $\Delta Y^\Omega(t)$ is bounded by a constant and apply (C.1b).

References

[1] P. C. Bressloff, Stochastic Processes in Cell Biology, Springer, 2014.

[2] A. Crudu, A. Debussche, O. Radulescu, Hybrid stochastic simplifications for multiscale gene networks, BMC systems biology 3 (1) (2009) 1.

[3] W. E, D. Liu, E. Vanden-Eijnden, Nested stochastic simulation algorithms for chemical kinetic systems with multiple time scales, Journal of Computational Physics 221 (1) (2007) 158–180. doi:10.1016/j.jcp.2006.06.019.

[4] T. Jahnke, M. Kreim, Error bound for piecewise deterministic processes modeling stochastic reaction systems, Multiscale Modeling & Simulation 10 (4) (2012) 1119–1147.

[5] A. Chevallier, S. Engblom, Pathwise error bounds in Multiscale variable splitting methods for spatial stochastic kinetics (2016). arXiv:1607.00805.

[6] T. G. Kurtz, The relationship between stochastic and deterministic models for chemical reactions, The Journal of Chemical Physics 57 (7) (1972) 2976–2978.

[7] D. F. Anderson, T. G. Kurtz, Stochastic Analysis of Biochemical Systems, in: Stochastics in Biological Systems, Vol. 1, Springer International Publishing, 2015, p. 90. doi:10.1007/978-3-319-16895-1.
URL http://www.springer.com/us/book/9783319168944

[8] D. F. Anderson, M. Koyama, Computational complexity analysis for Monte Carlo approximations of classically scaled population processes, IMA Journal of Numerical Analysis 35 (4) (2015) 1757–1778.
[9] D. T. Gillespie, Exact stochastic simulation of coupled chemical reactions., J. Phys. Chem. 81 (1977) 2340–2361.
[10] D. T. Gillespie, Stochastic simulation of chemical kinetics, Ann. Rev. Phys. Chem. 58 (2007) 33–55.
[11] S. Zeiser, U. Franz, O. Wittich, V. Liebscher, Simulation of genetic networks modelled by piecewise deterministic markov processes, IET Syst. Biol. 2 (2008) 113–135.
[12] Y. Cao, D. T. Gillespie, L. R. Petzold, Efficient step size selection for the tau-leaping simulation method., J. Chem. Phys. 124 (2006) 044109.
[13] D. F. Anderson, A. Ganguly, T. G. Kurtz, Error analysis of tau-leap simulation methods, The Annals of Applied Probability (2011) 2226–2262.
[14] D. F. Anderson, D. J. Higham, Y. Sun, Complexity of multilevel Monte Carlo tau-leaping, SIAM Journal on Numerical Analysis 52 (6) (2014) 3106–3127.
[15] D. F. Anderson, D. J. Higham, Multilevel Monte Carlo for continuous time Markov chains, with applications in biochemical kinetics, Multiscale Modeling & Simulation 10 (1) (2012) 146–179.
[16] J. B. Goodman, K. K. Lin, Coupling control variates for Markov chain Monte Carlo, Journal of Computational Physics 228 (19) (2009) 7127–7136.
[17] M. B. Giles, Multilevel Monte Carlo methods, Acta Numerica 24 (2015) 259–328.
[18] A. Ganguly, D. Altintan, H. Koeppl, Jump-Diffusion Approximation of Stochastic Reaction Dynamics: Error bounds and Algorithms, arXiv (2014) 32.
URL http://arxiv.org/abs/1409.4303
[19] S. N. Ethier, T. G. Kurtz, Markov processes : characterization and convergence, Wiley, 1986.
[20] K. R. Sanft, H. G. Othmer, Constant-complexity stochastic simulation algorithm with optimal binning, The Journal of Chemical Physics 143 (7) (2015) 074108. doi:10.1063/1.4928635.
URL http://scitation.aip.org/content/aip/journal/jcp/143/7/10.1063/1.4928635
[21] E. Levien, P. C. Bressloff, A stochastic hybrid framework for obtaining statistics of many random walkers in a switching environment, Multiscale Modeling Simulation In press.
[22] P. G. Hufton, Y. T. Lin, T. Galla, A. J. McKane, Intrinsic noise in systems with switching environments, arXiv preprint arXiv:1512.00785.
[23] F. Malrieu, Some simple but challenging markov processes, Annales de la Faculté des Sciences de Toulouse. Mathématiques. Série 6 24 (2015) 857–883.
[24] C. W. Gardiner, Others, Handbook of stochastic methods, Vol. 4, Springer Berlin, 1985.
[25] A. Duncan, R. Erban, K. Zygalakis, Hybrid framework for the simulation of stochastic chemical kinetics.
[26] M. G. Riedler, Almost sure convergence of numerical approximations for piecewise deterministic markov processes, J. Comput. Appl. Math. 239 (2013) 50–71. doi:10.1016/j.cam.2012.09.021.
URL http://dx.doi.org/10.1016/j.cam.2012.09.021
[27] R. Veltz, A new twist for the simulation of hybrid systems using the true jump method, arXiv [math].
URL http://arxiv.org/abs/1504.06873
[28] J. M. Newby, Isolating intrinsic noise sources in a stochastic genetic switch, Physical Biology 9 (2) (2012) 026002.
[29] R. Karmakar, I. Bose, Graded and binary responses in stochastic gene expression., Phys. Biol. 1 (197-204).
[30] P. Thomas, N. Popović, R. Grima, Phenotypic switching in gene regulatory networks, Proceedings of the National Academy of Sciences of the United States of America 111 (19) (2014) 6994–6999. doi:10.1073/pnas.1400049111.
URL http://www.pnas.org/content/111/19/6994$\delimiter"026E30F$nhttp://www.ncbi.nlm.nih.gov/pubmed/24782538
[31] P. Brémaud, Markov chains: Gibbs fields, Monte Carlo simulation, and queues, Vol. 31, Springer Science & Business Media, 2013.
[32] D. F. Anderson, M. Koyama, An asymptotic relationship between coupling methods for stochastically modeled population processes, IMA Journal of Numerical Analysis 35 (4) (2015) 1757–1778.
[33] D. F. Anderson, B. Ermentrout, P. J. Thomas, Stochastic representations of ion channel kinetics and exact stochastic simulation of neuronal dynamics, Journal of computational neuroscience 38 (1) (2015) 67–82.
[34] S. Isaacson, The reaction-diffusion master equation as an asymptotic approximation of diffusion to a small target, SIAM Journal on Applied Mathematics 70 (1) (2009) 77–111. doi:10.1137/070705039.
URL http://epubs.siam.org/doi/abs/10.1137/070705039
[35] S. A. Isaacson, D. Isaacson, Reaction-diffusion master equation, diffusion-limited reactions, and singular potentials, Physical Review E - Statistical, Nonlinear, and Soft Matter Physics 80 (6).
[36] J. P. Hespanha, A. Singh, Stochastic models for chemically reacting systems using polynomial stochastic hybrid systems, International Journal of robust and nonlinear control 15 (15) (2005) 669–689.
[37] A. Singh, J. P. Hespanha, A derivative matching approach to moment closure for the stochastic logistic model, Bulletin of mathematical biology 69 (6) (2007) 1909–1925.